

AMENDMENTS TO THE CLAIMS

Claims 1-7 (canceled)

8. (previously presented) In an information retrieval system having a current query, and a search engine that is responsive to a query to retrieve documents generally relevant to the query from a document collection, a method for refining the current query, the method comprising:

receiving information regarding the relevancy of documents retrieved at least in part in response to the current query;

ranking the retrieved documents at least in part in accordance with the relevancy information;

forming at least one candidate query based at least in part on one or more of:

the rankings and,

analysis of locations of the retrieved documents in the latent semantic index vector space formed from at least the retrieved document;

applying at least one candidate query to the document collection;

ranking the documents retrieved in response to each applied candidate query in accordance with the received relevancy information;

comparing the ranking of documents retrieved in response to at least one applied candidate query and the ranking of documents retrieved in response to the current query with the received relevancy information; and

choosing the query, from among the current query and each candidate query, which produces the best ranking.

9. (currently amended) A method for testing a hypothesis for refining a strategy for ranking the results of a document retrieval operation, the method comprising:

indexing results of a document retrieval operation ~~in to~~into a latent semantic index vector space;

receiving information regarding the relevancy, with respect to information needs, of a first subset of the results;

ranking a second subset, ~~nonexclusive of the first subset,~~ of the results based on a current ranking strategy;

modifying the current ranking strategy to incorporate at least one modification based at least in part on:

location of the second subset in the vector space, and

a subset of the relevancy information;

ranking the second subset based on the modified ranking strategy; and

determining which of:

the ranking based on the current ranking strategy, and

the ranking based on the modified ranking strategy

better corresponds to the relevancy information.

10. (currently amended) In an information retrieval system a method for refining a current query, the method comprising:

ranking a first subset of documents retrieved in response to a current query with respect to relevancy of the ~~retrieved documents~~ first subset to an information need,
ranking a second subset of the documents retrieved in response to a candidate query with respect to relevancy of the ~~retrieved documents~~ second subset to the information need;

wherein ranking is a function of at least the context of the relevancy information and each document subset in a latent semantic index vector space; and

choosing as a refined query, the query which produces a ranking indicating that the retrieved data corresponding to the query more closely matches the relevancy information.

11. (currently amended) A method for forming a hypothesis for refining a query, the query addressable to a ~~data collection, the data collection containing documents, each document containing a plurality of terms,~~ document collection the method comprising:

indexing a document collection into a latent semantic vector space;

receiving information regarding the relevancy, with respect to information needs, of a subset of the ~~data~~ collection;

forming at least one hypothesis regarding modifications to a query based at least in part on ~~analysis of the context of the relevancy information and the~~ location of documents corresponding to the relevancy information in the vector space.

12. (new) A method for refining a query, the method comprising:
receiving information regarding the relevancy to an information need of documents
retrieved from a document collection at least in part in response to a current
query;
ranking the retrieved documents at least in part in accordance with the relevancy
information;
forming at least one candidate query based at least in part on at least one of:
the rankings and,
the locations of the retrieved documents in the latent semantic index vector space
formed from at least the retrieved document;
applying the at least one candidate query to the document collection;
ranking the documents retrieved in response to each applied candidate query at least in
part in accordance with the received relevancy information;
comparing the ranking of documents retrieved in response to at least one applied
candidate query and the ranking of documents retrieved in response to the current
query with the received relevancy information; and
choosing the query, from among the compared queries, which produces the best ranking.